MINERvA neutrino detector response measured with test beam data

MINERvA Collaboration

MINERvA Addresses placeholder in template

Abstract

The MINERvA collaboration operated a miniature, MINERvA-like solid scintillator sampling calorimeter in a hadron test beam at the Fermilab Test Beam Facility. This article reports measurements with samples of protons, pions, and electrons from 0.35 to 2.0 GeV/c momentum. From these data, calibrations of Birks' constant, calorimetry, and a test of proton tracking efficiency are obtained. These results are used to tune the MINERvA detector simulation and evaluate systematic uncertainties in support of the MINERvA neutrino cross section measurement program.

Keywords: `elsarticle.cls`, LaTeX, Elsevier, template

2010 MSC: 00-01, 99-00

Note to MINERvA reviewers

This list will allow for quick location of additional information and supporting plots from our tech notes, and will be removed before being submitted to NIM.

Test beam beamline resolution and systematic errors
Devan and Gran, TN017, docdb:8547

Test beam detector calibrations, Gran et al., TN018 docdb:8686

Pion reaction cross section, Higuera, TN033 docdb:9031

Birks' law parameter from MINERvA testbeam data, Gran, TN037 docdb:9131

Test beam pion calorimetry, Gran, TN045 docdb:9474

Electron calorimetry and the e/pi ratio, Bergan, docdb:9929

Test beam proton calorimetry, Devan, TN051 docdb:9986

Short track reconstruction efficiencies using test beam protons,
McGivern, TN048, docdb:10367

1. Introduction and test beam goals

The MINERvA experiment[1] is designed to make measurements of neutrino nucleus cross sections with high precision. An important part of these [2, 3, 4, 5, 6] and future cross section measurements is the estimate of the energy of one or more hadrons exiting the nucleus. These are moderate energy recoil nucleons with energy from hundreds of MeV to a few GeV, especially protons and pions from inelastic production, plus softer nucleons and nuclear fragments. This goal of the test beam experiment is to validate the Monte Carlo simulation of the detector response to these particles. Significant results presented in this paper include a measurement of the Birks' law parameter, constraints on proton, pion, and electron calorimetry, and a study of tracking efficiency for protons.

The detector used to take these data is identical in most respects to the MINERvA detector installed in the NuMI neutrino beam at Fermilab, but in miniature. These test beam data are the first from a new hadron beamline at the Fermilab Test Beam Facility (FTBF) for a data run in summer 2010 as Fermilab Test Beam Experiment T977. There are differences between the two detectors that mitigate special aspects of the beam environment in FTBF and allow for a data set better focused on the Birks and calorimetry results.

The energy range covered by these data is well matched to the energy range of protons and pions and electromagnetic showers in the 2010 to 2012 MINERvA low energy neutrino and anti-neutrino data. This is especially true for the reactions from quasi-elastic scattering through Δ and other resonance production. Measuring differential cross sections for these exclusive final states is a pillar of the MINERvA neutrino physics program. These energies also cover the lower part of the range expected for hadrons produced in neutrino deep inelastic scattering.

2. Fermilab Test Beam Facility tertiary hadron beam

This beam is available because of a partnership between the MINERvA experiment and the Fermilab Test Beam Facility. It is produced from secondary

16 GeV pions colliding with a copper target to create tertiary particle production. All species exit a collimator at an angle of 16 ± 1 degrees. The species and momentum is tagged particle by particle using time of flight (TOF) and position measurements from four wire chambers (WC). A diagram showing the geometry and coordinate system is Fig. 1 with the beam going left in this top view and encountering the target and collimator first. The two magnets each carry 100

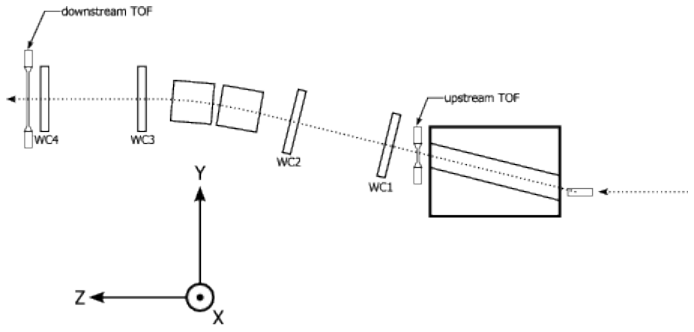

Figure 1: Diagram of the beamline built for this experiment, viewed from above with the beam going from right to left.

A of current and the polarity can be reversed. The typical field integral is 38.3 ± 1.5 Tesla cm with the latter number describing variations that encompass 90% of selected events. The detector sits off the left side of the diagram. It sees a range of incident particles with low momentum at low Y coordinates and normal incidence and high momentum at high Y coordinates and with angles as high as 10 degrees with respect to the detector Z axis.

With the large detector aperture (described later) and our chosen beam tune, the beam delivers a broad distribution of protons and pions from 0.35 to 3.0 GeV/c momentum. These particles are species and momentum tagged on a particle by particle basis. The usable momentum range for these analyses is 0.35 to 2.0 GeV/c which provides roughly ten thousand particles each of proton, π^+, and π^-, enough that these measurements are systematics dominated. The electron content of the beam is small and limited to low momentum, but has enough events for analysis. In addition, there is a 5% component of kaons, plus

smaller components of deuterons and alpha particles which are not part of the results presented here.

The resulting distribution is illustrated in Fig. 2 after quality cuts for well-measured particles. The pion, kaon, proton, and deuteron/alpha components are clearly seen, along with low momentum electrons near 20 ns. There is also an accidental background near 40 ns when another particle coincidentally passes through the TOF. This happens because the Fermilab Main Injector Accelerator supplying the beam has a 53 MHz time structure where pions striking the copper target are likely separated by an integer times 19 ns .

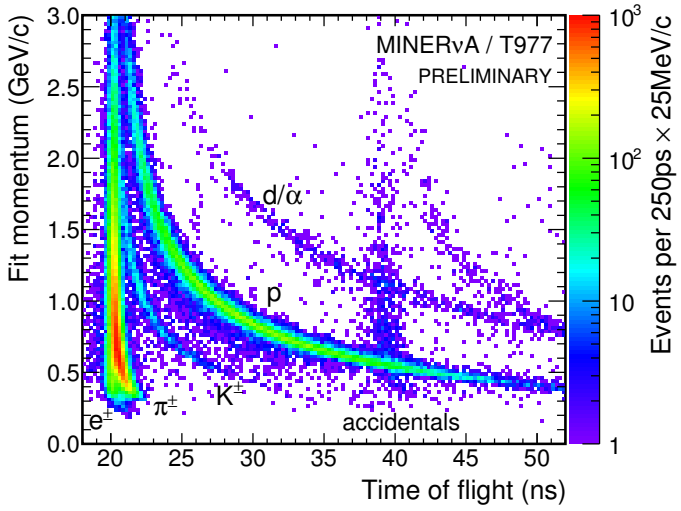

Figure 2: The measured momentum and time of flight used to separate different particle species and backgrounds.

The momentum estimate uses a detailed map of the magnetic field calculated using finite element analysis software from the specifications for the two dipole magnet coils and steel and the position survey of their placement relative to each other. The central value of the magnetic field from the calculation is adjusted down to match the actual field of the magnet from in-situ measurements with better than 0.5% precision. Longitude and transverse measurements of the field were taken with the magnets installed in their final positions. These

measurements are well described by the calculated field and how accurately they are modeled is one basis for estimating the uncertainty in the momentum. The other uncertainty comes from the accuracy of the position survey of the four wire chambers. In total, the systematic bias uncertainty on the momentum assignment ranges from 1% at low energy to 2% at higher energy.

The iterative momentum fit uses a stepper within the non-uniform calculated field to estimate the field integral, then a Kalman Filter technique is used to obtain the momentum and its uncertainty for each trigger. The momentum resolution is evaluated particle-by-particle and is 2.5% for pions and ranges from 5% to 3% for protons. It is driven by multiple scattering and non-uniform magnetic field effects at low momenta and by the wire pitch and beamline length at high momenta. The resolution of the momentum estimate is modeled accurately enough and is not a limiting factor for these analyses. The 200 ps resolution of the time of flight system limits the ability to separate electrons from pions. This TOF resolution is a source of uncertainty for those calorimetry results.

3. MINERvA test beam detector and calibration

The detector exposed to this FTBF beam (hereafter called the test beam detector) is a miniature version of the full MINERvA detector installed in the NuMI neutrino beam [1]. It is made of forty rectangular planes of scintillator with length 1.07 meters and 63 strips wide. This contrasts to the full MINERvA detector which has a hexagon shape made of 124 planes of 127 strips in the central tracker region followed by another 20 planes each of ECAL and HCAL with Pb and Fe interleaved respectively. Both detectors share the same three-view UXVX sequence of planes with U and V rotated ± 60 degrees relative to the X plane that defines a vertical coordinate system. This allows for dual stereo reconstruction of multiple tracks for the full MINERvA and very good reconstruction of single tracks in the test beam detector.

The signal chain from scintillator to wavelength shifting (WLS) fiber to PMT to digitization is almost identical. The exception is the test beam detector has

no clear fiber optical cables; the WLS fiber connects directly to the PMT a half-meter out of the plane. The effect of smaller scintillator planes and no clear fiber is that the test beam detector has about 50% higher light yield, and corresponding better resolution for some kinds of measurements, compared to the full MINERvA.

Unlike the MINERvA detector which reads out every plane on the same side, the test beam detector alternates readout in groups of four planes, one UXVX set rotated 180 degrees. Mechanically this allows the planes to be placed closer together, with an air gap only slightly larger than the MINERvA detector. Because the beam bend magnets steer different momentum particles to different portions of the detector (and at different angles) there is a correlation between the geometry and the position dependent optical attenuation of the readout. Alternating the readout mitigates a few-percent momentum dependent uncertainty, making the effect negligible.

The detector energy scale is calibrated using the same strategy as the MINERvA detector installed in the NuMI neutrino beam, and is described in [1]. An initial estimate for photoelectrons is obtained for each strip using standard pedestal and gain measurements. The intrinsic differences between strips are analyzed to produce a correction factor to make the response from strip to strip uniform. As a side effect, these muons also gives geometric plane position corrections. Then the absolute energy scale is done using a muon equivalent unit (MEU) technique. First this ensures the peak number of photoelectrons at the PMT is the same in the data and simulation. Then the simulated geometry and Geant4 energy loss is used to give the absolute energy scale.

The absolute response calibration is carried out with broad spectrum cosmic ray muons (and a simulated spectrum) rather than momentum-analyzed muons from the NuMI beam. These calibrations do not include energy that appears off the muon track due to cross talk, a feature treated separately in the analyses described in this paper. The calibration uses the peak of the dE/plane response for muons, and depends little on muon δ-ray and bremsstrahlung production in the tail of that distribution.

Secondly, temperature dependence is more important than it is in the NuMI hall. The detector hall warmed during the day and cooled at night, so the overnight cosmic muon sample spans the same range of temperature as the daytime hadron sample. The detector response is corrected for that temperature dependence and a residual uncertainty is included with the systematic errors.

Unlike the MINERvA detector, the test beam detector's removable absorber planes allow for exposures of two configurations. One has 20 planes with 1.99 mm thick Pb absorber (ECAL) and 20 planes with 26.0 mm thick Fe absorber (HCAL). The absorber is interleaved by placing one absorber upstream of each scintillator plane. The other has 20 planes with no absorber (tracker) and 20 planes of ECAL. For compactness, this document will refer to these configurations as EH and TE, respectively. This replicates the main downstream regions of the MINERvA detector, which has 124 planes of tracker followed by 20 planes of ECAL and 20 planes of HCAL.

4. Data sample and simulation

For the EH configuration, the following figures show the energy spectra for π^+, π^-, and protons.

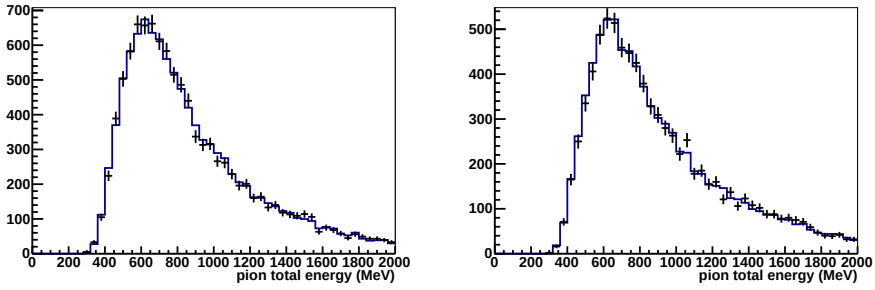

Figure 3: Measured spectra for π^-, π^+, and proton samples, after selection. The blue histogram is taken from a MC simulation that was seeded with measured particle momenta and trajectories from the data. [ADD PROTON EH PLOT, maybe only show one pion plot as an example.]

In this analysis, the data are compared to a full, high statistics simulation. The spectrum for the simulation are the actual data particles' position and momentum measured at the third wire chamber (WC3), with momentum and angle smeared according to the estimated resolution on a particle by particle basis. The simulation then propagates particles through the material of WC3, WC4, the downstream TOF, the cosmic muon trigger scintillator, and the air, then finally into the detector. Using the estimated resolution for each particle and a Gaussian random smearing allows us to use the same data particles multiple times and generate large MC samples, typically 20 to 40 times. We observe that the MC reproduces the data for the special situations where the beamline resolution dominates, such as for protons that stop at the end of their expected range, validating that the beamline characteristics are well simulated.

The MC does not simulate any beamline induced background effects, neither from particles that are exactly in-time (from the same parent 16 GeV pion hitting the target) nor secondaries from another pion in an nearby 19 ns slot in the Fermilab Main Injector 53 MHz accelerator structure. Because we collect data for 16 microseconds around each trigger, and because some incident particles should spatially leave much of the detector quiet, the data itself contains a record of the average beam-induced background around valid triggers. These estimates were validated, and some cleaning selections were developed using the revolutionary web-based MINERvA event display [7], in many cases with the help of undergraduate research assistants. These selections reduce both beam-induced backgrounds and also eliminate triggered particles that scattered substantially in the beamline before reaching the detector. The selections to reduce these unwanted events are applied to both the data and simulated samples. Finally, we estimate and make a statistical subtraction of the remaining background and evaluate an uncertainty for each analysis.

The basis of the simulation uses Geant4 version 9.4p2 [8, 9] and our best description of the detector geometry and material [1]. The scintillator plane is made of 1.801 g/cm^2 of plastic scintillator, WLS fiber, and a co-extruded TiO_2 reflective coating. Added to this is another 0.226 g/cm^2 of epoxy and Lexan.

The scintillator strips were made at the same facilities immediately following the production of MINERvA planes, and the manufacturing modifications for smaller planes make negligible difference. In the ECAL portion of the detector there are planes of Pb with thickness 2.30 g/cm^2 and in the HCAL version there is 20.4 g/cm^2 material that is 99% Fe and 1% Mn. The Pb and Fe absorber are similar to the MINERvA detector, but we use the as-measured test beam detector quantities in the simulation and to evaluate material assay uncertainties.

Almost all aspects of the detector response are simulated using details constrained by calibration data and bench tests, including the Birks' law parameter measured from these data, described below. A few features are not simulated, of which PMT after-pulsing and PMT non-linearity are the only significant omissions. All uncertainties in this category that significantly affect the analysis are described later.

5. Birks' Law parameter

With a well calibrated beam and detector, the test beam data are used to obtain a measurement of the Birks' Law parameter [10, 11] for MINERvA scintillator. The scintillator is made of polystyrene doped with 1% PPO and 0.03% POPOP by weight [1]. Birks' Law describes the quenching effect where the conversion of energy to scintillation photons is suppressed for high, localized dE/dx. This is especially important at the end of proton tracks in neutrino analyses, which why we need this calibration. Subsequent measurements in this paper rely on the determination of this parameter. The correctly adjusted scintillator response and its uncertainty play a role in the calorimetry measurements. Stopping test beam protons are the means to obtain the required calibration. This measurement is done with protons that appear to stop in the upstream 20 tracker planes of the TE detector configuration.

A sample of protons that stop at the end of their expected range is selected from the larger set of good proton events. Protons that appear to stop between

planes 9 and 19 inclusive are checked that their range is consistent with their incoming energy. It is not done using a simulation; for events that stop in a particular plane, the end-of-range protons form a peaked distribution at the appropriate kinetic energy while protons that interact are a tail extending to higher energy; we select the former.

Once selected, the incoming energy of the protons (according to the beamline measurement) is not used for analysis, the dE/dx profile is built backward from the observed end of each proton track. This is done separately for the data and MC. Compared to analyzing as a function of incoming energy, this reduces sensitivity to uncertainties or mismodeling of the incoming proton momentum or material assay. On this latter point, we observe that the proton range is very well modeled by the simulation. The simulated protons stop 1.1% earlier than the data, which is a smaller discrepancy than the 1% beamline momentum and 1.4% material assay uncertainties. This cross-check uses a Gaussian fit to find the end-of-range peak used to select the Birks' sample, and is shown in Fig. 4. Stopping protons are such a high resolution sample, the widths of those Gaussian fits (10 to 15 MeV, not shown) are primarily driven by the beamline and multiple scattering resolutions, and are also well described by the simulation.

With this selected sample, the activity from each proton is analyzed as a function of how many planes it is from the last plane with activity. The last plane with activity is not always the plane with the most activity. The binned distribution of measured energy in each plane-from-the-end will be used in the fit, but to ease visualizing the trend, a Gaussian can be fit to the peak and plotted as in Fig. 5. Both the binned distribution and the fitted distribution are done for the data, the default simulation before fitting for a better Birks' parameter, and two alternate simulations with $\pm 30\%$ variation of the nominal parameter. The points in this visualization are from a Gaussian fit, which is a good approximation for most planes, but because the proton stops at different depths in the final plane, the distribution is not actually Gaussian. The trend clearly indicates the original estimate of Birks' parameter needs to be shifted to

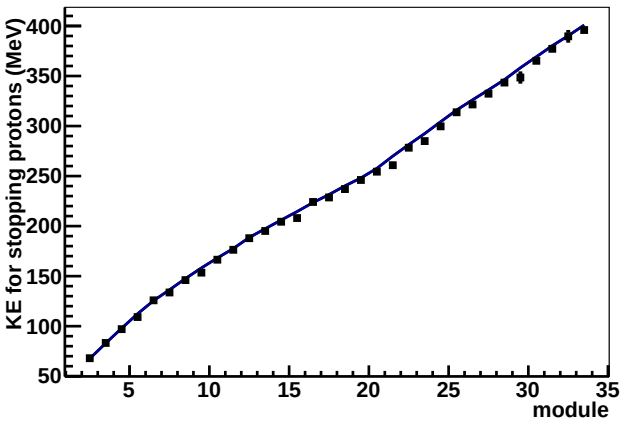

Figure 4: Kinetic energy from a Gaussian fit to the peak energy for protons that appear to stop in each TE plane (module). The MC stops 1.1% earlier than the data, a discrepancy which is smaller than the beamline momentum and material assay uncertainties. Most error bars are less than 1% and are too small to see.

produce less quenching.

The Birks' parameter is obtained by using the default simulation with a candidate Birks' parameter and two other simulations with separate ± shifted Birks' parameters, and interpolating to get good agreement with the data. The data used for the fit are the inputs to Fig. 5, the binned profile of dE/plane for each plane back from the end of the track. The fit is limited to bins of dE/plane where the statistics are good. When each bin of dE/plane and each plane-from-the-end are combined, there are 123 bins total in the fit. In the fit, the overall energy scale is an unconstrained parameter, which simultaneously accounts for both the energy scale uncertainty and the correlation between the calibrated energy scale and Birks' parameter. Also, an amount of dE smearing that accounts for unsimulated calibration effects was allowed to vary, though it yields the same 5.5% result [1] as found for the full MINERvA. Thus the fit is done using a parameter scan in this 3D parameter space of Birks' parameter, energy scale, and smearing of reconstructed energy deposits. The procedure

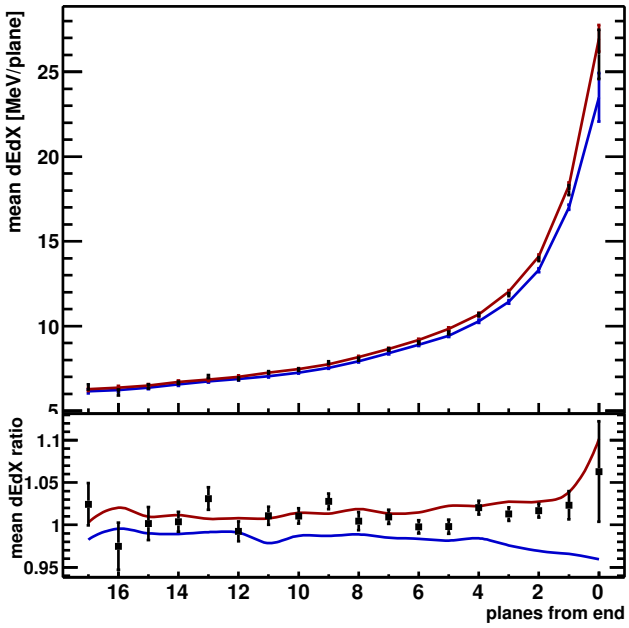

Figure 5: The measured energy deposit per plane for data compared to the simulation with the before-fit Birks' parameter of 0.133 mm/MeV and the original -30% (red) and +30% (blue) uncertainty on this parameter's initial value.

is iterated with a new MC built from the new parameter value and smaller ± shifted values; the iteration reliably converges to the final result.

The best value for the Birks' parameter is 0.0905 ± 0.015 mm/MeV. This value is near the -1σ limit of the ab-initio estimate used by MINERvA for analyses through 2014, confirming we use suitable Birks' effect uncertainties in prior publications. Future simulations using the new value have half the prior uncertainty. The best fit describes the data well, yielding a χ^2 of 124 for 120 degrees of freedom.

This is a systematics dominated measurement. The main systematic for the Birks' parameter measurement comes from the the correlation with the energy scale. Uncertainties in the material assay and the variations among the limited set of physical planes being tested contribute additional smaller amounts to the

total.

Because this parameter is obtained by matching the MC simulation to data, it might be considered an effective parameter. In addition to energy scale correlations and Birks' quenching, it is accounting for the accuracy of the Geant4 energy loss simulation and our choice to use the default (adaptive) Geant4 step size, in addition to the actual Birks' quenching effect. A more coarse aggregation of Geant4 steps at the scale of one scintillator bar yields an increase in the response of about 4% in the last plane and a slightly better $\chi^2 = 118$. The typical dx has increased, so dE/dx has decreased, so there is less Birks' suppression. Such an extreme effect would cause a bias in the fit Birks' parameter of 0.008, about half the total uncertainty. However, this particular measurement is specifically matched to the Geant4 and hit aggregator settings that are used by the full MINERvA simulation as of late 2014, so this is a cross check and not an uncertainty.

Another detector response parameter that has systematic effects, especially on calorimetry, is the nonlinear response of the photo-multiplier tubes due to saturation effects in the dynode current. This non-linearity sets in for high instantaneous current at the anode, and so is a function of charge measured by the front end board's digitization module. As of this writing, MINERvA does not have an in-situ measurement under circumstances that are equivalent to the light propagating in our scintillator bars and WLS fiber. Instead, we have a reference non-linearity curve obtained from bench tests. The dE/plane profile in Fig. 5 is distorted by nonlinearity in ways different from either Birks' parameter or energy scale, so we investigate the size of possible nonlinearity. Applying nonlinearity of 20% of the reference degrades the χ^2 by one unit, with a correlated shift in Birks' parameter. Since applying some nonlinearity does not give an improvement in the χ^2, the best fit value is effectively none. Thus at 20 MeV per plane, we do not have sensitivity to nonlinearity effects with these data. We use this 20% to get an uncertainty for the Birks' parameter measurement. A larger uncertainty of 50% of the reference is used later for calorimetry to account for non-linearity effects when higher energies are deposited in a single strip.

These results can be compared to other values for the Birks' quenching parameter. A recent review of of the properties of many materials including polystyrene is available in [12] with references and one additional later measurement [13]. These measurements are focused on heavily ionizing nuclear fragments and alpha particles which are important in dark matter and double beta decay experiments as well as nuclear fission studies. The technique is conceptually similar to using the end of a proton track but potentially more sensitive due to the enhanced ionization and granularity of the data. The analysis of [12] obtains a value of 0.0090 g / cm^2 MeV (with no uncertainty given) for polystyrene based scintillator. Using the 1.06 g/cm^3 density of polystyrene quoted in that analysis, this converts to 0.085 mm/MeV. This value and the Birks' parameter result above for our scintillator formulation and density of 1.043 g/cm^3 are nearly identical. There are a multitude of additional values for Birks' Law parameter used by neutrino experiments for the variety of available plastic scintillator and other materials. The parameter value is expected to depend primarily on material formulation.

6. Proton calorimetry

This test beam experiment is designed to constrain the uncertainty on the single particle calorimetric response to protons and pions. For low multiplicity neutrino events we reconstruct the hadron response particle by particle using range, calorimetry, or a combination of the two. For high multiplicity hadron systems from neutrino events, the total energy of the hadronic recoil system (everything but the outgoing charged lepton) is calorimetrically reconstructed. In both cases, when the hadron(s) interact energy is lost to unbinding nucleons and to neutral particles. An estimate of this missing energy is used to correct the observed response and obtain an unbiased estimator for the hadron system. In all cases, a major ingredient is the MC prediction for the single particle response, which is constrained with these data.

The hadron event is reconstructed by summing the calibrated energy mea-

sured in the scintillator. The standard tracking algorithm is applied to each event. If a track segment is found, the 3D location of hits on the track are known and used to make a correction for attenuation in the scintillator strip to the point where the particle passed. For all hits not on tracks, the attenuation estimate is made to the center of the strip. Then a correction for the passive material fraction for each plane is applied; a factor of 1.255 in the tracker, 2.077 in the ECAL, and 10.727 in the HCAL. The muon equivalent unit (MEU) technique does not include cross talk when setting the energy scale. Because cross talk is proportional to the total of the energy deposits, the measured cross talk fraction of 4.2% is subtracted from both data and MC. The activity recorded over the 150 ns digitization time is used, unlike the typical MINERvA neutrino analysis which uses a window from -20 ns to +35 ns around the peak in the cluster timing distribution. The latter technique is tested as a cross check. Activity later than 150 ns from low energy neutrons and decay electrons is not included. The former is predicted to amount to a few percent of the available energy and appears in the detector over several microseconds.

For the proton calorimetry analysis, the beamline induced backgrounds are reduced using additional selections. In the EH detector configuration for most energies, it is reduced by using activity in the last four planes as a veto. For the lowest proton energies, the back half of the detector is not included calorimetrically at all.

The corrected estimate for the energy is compared to the available energy, which is just the kinetic energy for the proton. This fractional response is calculated event by event. Then the events are binned by energy, from which we compute the mean and RMS for each bin. The results for the mean are plotted in Fig. 6. The error band on the MC represents the total systematic uncertainty.

The proton response has several features in this energy range. At low energy, the probability for a proton interaction is low. The result is that most of the energy is measured for most of the protons, and the distribution of response is approximately Gaussian around this mean. At 0.3 GeV, the protons

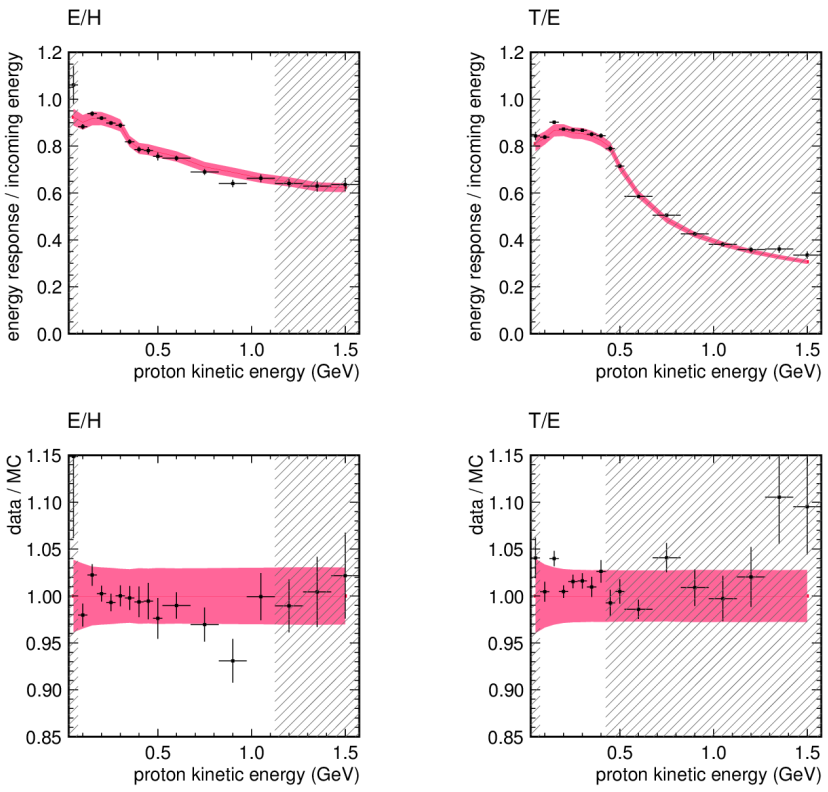

Figure 6: Proton response for EH left and TE right. The bottom plots show the ratio data/MC. The hatched region shows the response in for energies where containment is too poor for calorimetry. See text for discussion.

simultaneously begin to produce Δ resonances when they interact in nuclei and also penetrate to the HCAL. The former leads to a drop in response because Δ production generically leads to lower response through neutral final states and unbinding of additional nucleons.

The MC tracks the proton response well over the entire range. This is shown in the ratio data/MC for the mean response in each energy bin. The MC has negligible statistical error, the systematic uncertainty is shown, and the data is shown with statistical uncertainties in the bottom plots of Fig. 6. There is mild evidence of systematic effects consistent with any combination of

several calibration uncertainties, described later in its own section. There is a special pion background at 0.15 GeV in the data because those protons take 19 ns to travel the beamline, the same as the timing of the 53 MHz accelerator RF structure. These protons in data and MC and their backgrounds in data are cut from the analysis already. This data point, and the data point at 0.9 GeV corresponds to no special features of the experimental setup, and have the character of a $\sim 3\sigma$ fluctuation.

The response at low energy for the TE detector is partly correlated to the tuning of Birks' constant, because up to 0.25 GeV they are the same proton events. However, the strip response energy scale does not come from the free parameter in the Birks' fit, which would make this correlation even more complete. Instead, the calibration muon sample was refit using the measured Birks' parameter to obtain the final strip energy calibration. Thus energy response offsets are correlated only with the Birks' parameter and its uncertainties.

In these figures we show the comparison of data and MC in a region at higher energy, which is shaded. At these energies in the TE configuration we are losing containment of charged particles produced in the hadronic interaction, and the calorimetric response no longer represents the kind of result we expect for the full MINERvA detector. Instead, these points demonstrate that the MC is still doing an adequate job describing the data.

In addition to the response, it is important for MINERvA neutrino analyses that we know the fluctuations in the response are well simulated. Many neutrino distributions are strongly peaked in reconstructed energy or some other kinematic quantity, and an error in resolution will flatten or sharpen the MC peak relative to the data, causing a bias in unfolded distributions and fit parameters.

The basic shape of the distribution of response particle by particle is well described, so it is adequate to use the RMS of the distribution to quantify the trend and the agreement, as shown in Fig. 7. The statistical uncertainty on the RMS is shown, no systematic uncertainty is quantitatively considered.

At all energies, the MC response has a lower RMS, more prominent at low energy. Though the deviation can be taken to be a conservative uncertainty

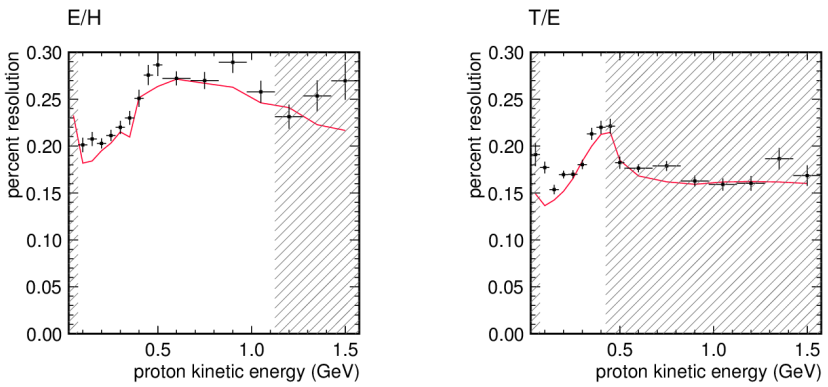

Figure 7: Calorimetric resolution in proton response for the EH configuration (left) and TE (right).

on the calorimetric resolution, the higher RMS in the data is partly from unsimulated beamline-induced activity, which are not expected to have the same magnitude effect for higher energy protons, the pion sample, or the same origin as events in the full MINERvA.

The resolution changes from 0.3 to 0.5 GeV because of two effects. This is the region where proton interactions start to produce Δ resonances, and also produces the decrease in response seen in the upper left plot of Fig. 6. In the EH configuration, this is also the energy range where protons are likely to reach the HCAL.

Of special interest is the resolution for the lowest energy protons which are contained in the tracker portion of the TE detector configuration. These correspond to the protons typically found at the vertex of a neutrino interaction from quasi-elastic and resonance production and include products of the intranuclear rescattering process. In this 0.05 to 0.2 GeV range a Gaussian fit to the peak gives a resolution of 18% for protons energetic enough to travel through more than one plane but are not energetic enough to excite Δ resonances in the nucleus. The resolution is driven mostly from fluctuations at the end of the range and (for data only) from beam-induced background activity. Above this point

Δ production becomes important, reducing the stopping component to about half the total. The distribution from which the RMS is computed picks up a low-side tail whose shape is well modeled by the MC.

7. Systematics for single particle response

The systematic uncertainties on the single particle response described here also apply to the pion measurement that with only a few differences, as noted later. These systematic errors are evaluated using a variety of methods. Some are done by varying the selection criteria for data and MC and evaluating inconsistencies in the changed response, others are done comparing a modified MC to the default MC, and some are done by dividing the data sample into halves according to natural variations in detector and run conditions. The largest sources of uncertainty are described in Table 1.

Source	TE p	EH p	EH π^+	EH π^-
Beam momentum	1.9%	1.9%	1.0-2.0%	1.0-2.0%
Birks' parameter	2.0 to 0.9	2.0 to 1.2	1.0	1.0
Correlated late activity	0.3	0.6	1.4	1.4
Temperature stability	1.0	1.0	1.0	1.0
Energy scale	0.6	0.6	0.6	0.6
PMT nonlinearity	0.7	0.7	0.9	0.9
Event selection	<0.2	<0.2	0.7	1.5
Cross talk	0.7	0.9	0.5	0.5
Beamline mass model	0.7	0.7	<0.2	<0.2
Total	3.3 to 2.7%	3.4 to 2.9%	2.6 to 3.4%	2.9 to 3.6%

Table 1: Percent systematic uncertainties on the single particle response for data vs. MC comparisons. Additional uncertainties on the energy scale and absorber material apply 2.0% equally to data and MC absolute response, as described in the text. The total range represents the evolution with energy from 0.1 to 0.4 GeV for TE protons, 0.1 to 1.0 GeV for EH protons, and 0.4 to 2.0 GeV for both pion samples.

7.1. Beam momentum

This uncertainty is intrinsic to the design of the beam and the estimate of the momentum of the incoming particle. An uncertainty here has the trivial effect of shifting the denominator of the fractional response. It comes from the wire chamber survey and the measurement and simulation of the magnetic field. Because it is an uncertainty on the momentum, it translates differently to uncertainties on the available particle energy for protons and pions. The lowest energy protons pick up an additional 0.7% uncertainty on the energy loss in the material of the beamline which is included in the total and the error band.

7.2. Birks' parameter

Even after producing a best fit Birks' parameter, the remaining improved uncertainty is still one of the largest contributions to the accuracy of the result. Because low energy protons almost always have a high dE/dx activity at the very end of the proton's range, and because that activity is a larger fraction of the total energy for low energy protons, that sample is most affected by this uncertainty.

7.3. Correlated late activity

Some systematics are revealed by varying event selection cuts. Proton fractional response, and especially pion fractional response changes when additional substantial activity is reconstructed within 800 ns following the triggered event. The response in the MC, which has neither beamline-induced backgrounds nor PMT afterpulsing simulated, rises because of the correlation with neutrons, electrons from π to μ to e decay, and other delayed activity. These follow systematically pions with low fractional energy response. The response for the data falls slightly and ends about 1% below where the MC predicts. Particles removed with this cut in the data due to unrelated beamline activity should be uncorrelated with the energy of the triggered event. Finally, data particles with large shower activity generate more afterpulsing and neutrons and are more likely to have activity reconstructed in this 800 ns window. If it was primarily

afterpulsing, this would not be a systematic, but an investigation did not confirm that hypothesis. This could be a Geant4 modeling effect, but we have not ruled out an experimental systematic, so this is included in the uncertainty.

7.4. Temperature stability

The response of the detector to cosmic ray muons for the data is calibrated against the measured temperature in the experimental hall as a function of time. This accounts for the change over the course of the day and from day to day during the run. This gives a correction which is then applied to energy deposits in the beam data. The simulation has no temperature dependence.

7.5. Energy scale

The absolute energy scale is dominated by the material model for the scintillator planes and affects both data and MC. The calibration procedure uses a comparison of simulated cosmic ray muons to measured muons, so by construction the data/MC relative energy scale is very well constrained. The largest contribution to this relative uncertainty comes from observations of discrepancies between the TE and EH data sets, and not from the design of the method or the statistics of the data samples. Within each subsample, there is no discernable time dependent trend that would extrapolate between these two detector configuration. The uncertainty listed here is taken to be half the discrepancy seen in the muon calibrations between the TE and EH data sets.

7.6. PMT nonlinearity

The uncertainty is evaluated by applying half the magnitude of the nonlinearity reference curve to recalculate the reconstructed energy of MC on a strip by strip basis. It has a large effect for rare high activity strips, but for hadronic tracks and showers at these low energies the overall effect is modest. This effect is one way, there is no PMT nonlinearity in the simulation, so it serves only to move the simulated energy lower.

7.7. Event selection

The proton sample does not pick up a significant uncertainty because it cuts around the pion background for KE = 0.15 GeV. The pion sample selection intrinsically allows in an electron and kaon background. Variations in those selections yield a 0.7% uncertainty for π^+ and twice the uncertainty for π^-.

7.8. Cross talk

A measurement of the cross talk in the cosmic muon calibration finds it contributes an average of $4.2 \pm 0.5\%$ to the energy in the detector, and the amount in the MC is tuned to reproduce this. Because the energy calibration of the detector specifically does not include cross talk, this amount is subtracted from the total, and the remaining 0.5% uncertainty contributes to the calorimetric uncertainty between data and MC. Analysis of neutrino data also has cross talk in the simulation tuned to the data, and uses multiple techniques to deal with cross talk, including thresholds, topological identification, and subtraction depending on the analysis.

7.9. Absolute energy scale

There are additional effects which apply equally to both data and MC absolute energy scale. These enhance the absolute uncertainty beyond to the relative energy scale uncertainties. The most important come from the material model for the scintillator planes and also the lead and steel absorber. They affect both the calibration of the energy deposits in the detector as well as how deep the hadronic activity propagates into the detector. They add an additional 2% in quadrature to the quantities in Table 1.

7.10. Geant4 step size

The simulation is affected by a number of different Geant4 settings, including some that are unrelated to the hadronic physics model. An interesting one is the fidelity obtained with the default Geant4 adaptive step size algorithm for

stepping particles through the simulated detector, subject to certain user specified maximum steps. Purposely making the maximum step size 0.05 mm allows the adaptive algorithm to still choose smaller steps near material boundaries but never larger steps. This change results in a reduced MC response of 1% for pions and no effect for 0.5 GeV/c electrons. This is consistent with triggering an enhanced Birks' effect because then the simulation produces more highly quenched energy deposits. The baseline simulation here uses essentially the default settings, the same as used for the full MINERvA, so all the calibrations and measurements are done with a consistent set of parameters, and there is no uncertainty to assign.

8. Pion calorimetry

There are pion samples with two polarities. The EH π^+ sample was obtained concurrently with the proton sample while the π^- sample was from the data set taken the previous week. After these data were taken, the detector configuration was changed to the TE configuration, but unlike for protons, containment in the TE is not adequate for a pion calorimetry measurement. Another difference is that the lowest beam momenta available cause the lowest pion energy for this analysis to be 0.35 GeV, just above the Δ production peak. The mean free path would be around 30 ECAL planes, so very few pions stop at the end of their range in the detector, but many reach the HCAL before interacting.

The event selection and energy measurement proceed similar to the proton case, correcting the observed energy for passive material, cross talk, and the last-four-plane veto. The denominator for the response for pions is taken to be the total energy. For pions there is a potential background at low energy from electron contamination and at high energy from kaons which is neither simulated nor subtracted. Using variations of the selection, this is evaluated to be a small uncertainty.

The remaining background has been estimated two ways, by inspecting activity 30 ns earlier than the triggered particle and for the lowest energy proton

sample by inspecting activity deep in the detector where there should be negligible activity. When extrapolating these methods in time and space, they both yield the same 4 MeV per event on average. For the mean response, this is simply subtracted. At higher energy, the use of the veto leads to another bias of about 1% estimated using the MC, because real hadron interactions put energy into those planes. This bias is repaired with a MC-based energy dependent correction. This procedure is different than is done for protons, but leads to the same negligible uncertainty.

The MC describes the response to pions well, but not perfectly, and is shown in Fig. 8. The statistical uncertainty on the data is shown while the same for the MC is negligible, and all systematic uncertainties (with their energy dependence) from Table 1 are incorporated into the MC error band. One way to summarize the agreement is that the MC models the single particle response to within 4% up to 1.0 GeV, and 3% up to 2.0 GeV.

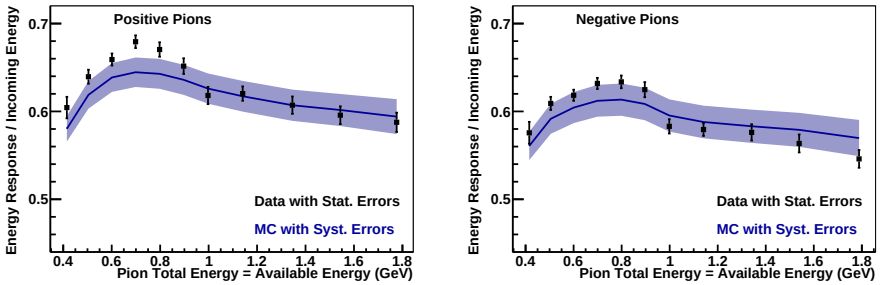

Figure 8: Calorimetric response for positive (left) and negative (right) pions. The errors on the data are statistical only, while the error band on the MC represents the systematic uncertainties on comparisons between data and MC. A larger uncertainty of up to 4.2% (not shown) applies to the absolute response scale both data and MC.

The MC error band is near or touching the data points' statistical error bars. This agreement is adequate for MINERvA's neutrino program, sets the uncertainty, and no correction is necessary. However, the MC does not accurately model a change in behavior that starts at 0.9 GeV. This change correlates with a mild inflection point predicted for the MC response. It also correlates

with a change in modeling the fraction of events with negligible energy in the HCAL, the MC underpredicts this starting at 0.9 GeV. The onset and the magnitude of the discrepancy are the same for both π^+ and π^-, equivalent to a -5% change from low to high energy relative to the MC. The experimental systematic uncertainties permit some discrepancy at higher energy relative to low, about equally from the beamline systematics, species selection, and beamline-induced backgrounds. When evaluated in quadrature, these could produce a $\pm 1.8\%$ relative change over this energy range, less than half what is observed. None of these systematics would naturally produce a change over a short 0.2 GeV energy range near 0.9 GeV. Instead of taking an overall uncertainty in the response, a MINERvA neutrino analysis sensitive to this will propagate the trend.

In principle, these data are a test of our ability to model the detector itself as well as the pion energy loss and reaction processes such as inelastic, absorption, charge exchange, and elastic scattering. We have investigated the sensitivity to model uncertainties using the Bertini Cascade model [14] within Geant4, including consideration of pion cross section data [15, 16]. However, calorimetry is more sensitive to the total available energy than it is to differences in outcome for any of these individual fates. Trial 30% modifications to the relative mix of fates have at most a 0.5% effect on calorimetry. Increasing the probability to interact (inelastic or elastic with at least 10 MeV energy transfer) before reaching the HCAL enhances the response. By this definition of interaction, the mean free path in the ECAL is about 30 planes; lowering it by about 20% increases the calorimetric response by 1.5%.

Uncertainties on the models in Geant4 in principle could be energy dependent, allowing a tuned model to better describe the overall average response or separately the anomalous trend with energy. Investigation of the trend reveals a correlation with the fraction of events that have negligible energy in the HCAL, the MC does not follow the data and underestimates this fraction starting at 0.9 GeV. This is also a predicted effect of lowering the mean free path, but detailed comparisons of the lateral and longitude distribution of energy is beyond the scope of this result. Simple followup investigations do not corroborate this hy-

pothesis, further analysis may be possible, including followup data being taken now. Though no experimental systematic by itself is expected to have these features, we do not rule out an artifact of the experimental setup or analysis.

The π^+/π^- ratio cancels a number of common systematics, including the trends described in the preceding paragraphs. The MC predicts that π^+ yield a 4.8% higher response than π^-, the measured ratio is 6.2%, consistent with flat across this energy spectrum for both. The statistical uncertainty in the data ratio is only 0.5% averaged over these energies. The 0.6% uncertainty in the energy scale comes primarily from an unknown time or detector configuration dependent effect, which should conservatively be applied to this ratio. There is no evidence for either an intensity effect (the π^+ data was at higher intensity) or an operational effect due to time or polarity in the beamline, nor a temperature effect. These are judged by comparing two halves of each data configuration further split along these operational parameters, though these tests are themselves afflicted by 0.7% statistical uncertainty. This discrepancy is at two standard deviations, and it can be used as a conservative uncertainty on the ratio.

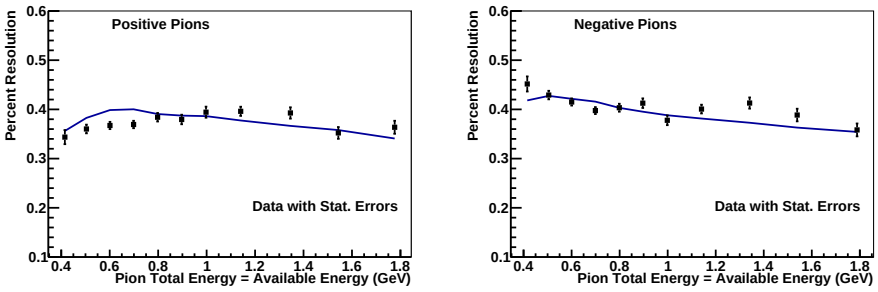

Figure 9: Resolution of the calorimetric response for positive (left) and negative (right) pions. The statistical error on the RMS is shown for the data, no systematic uncertainties are included.

As with the proton case, Fig. 9 shows the pion fractional response resolution is adequately modeled. The beam-induced backgrounds are a much smaller fraction of the total energy than for low energy protons, the Δ production peak

happens at energies below these, and a large fraction of the events reach the HCAL, so there is none of the structure seen in the proton case.

9. Electron calorimetry

The electron sample is limited to energies in a range from 0.4 to 0.5 GeV. The production of electrons is intrinsically lower in energy and fewer than pions, and the TOF resolution prevents good identification of the few that are at higher energies. Using the EH detector configuration, these electrons begin and complete >95% their electromagnetic shower in the ECAL portion of the detector, and the detector response in this special case can be isolated. The TE detector configuration is similar: the electron propagates through the tracker but does not shower extensively until the ECAL.

The electron sample is separated from the pion sample using a combination of topological and time-of-flight selections. Events that resemble late-interacting pions because they are tracked into the HCAL or because they have a substantial fraction of energy in the back half of the detector are rejected. Further, the number of strips recording activity is systematically more for electrons, and the variance in energy per plane for EM showers is much higher than for interacting pions. Using the MC, we estimate the efficiency for selecting electrons (pions) to be 61% (5%) for the TE and 73% (8%) for the EH configuration. The pion and electron peaks separate in time of flight by at least 0.7 ns at 0.5 GeV, well within the 0.2 ns resolution of the TOF. Extrapolating the pion distribution just above the TOF cut into the selected electron region in data yields an estimate of one pion background in 50 electron events. An eye-scan of the resulting events with the web-based MINERvA event display [7] confirms there is one obvious background event which is removed, leaving 49 events total.

The data and MC for electrons and positrons for the EH configuration were combined into one sample. The resulting sample is analyzed similarly as previously described for protons and pions. After correcting for passive material, cross-talk, and beamline-induced background activity, the response ratio is ob-

tained for every event. The electron response fraction is found to be 0.763 ± 0.013 (statistical) in data and 0.740 ± 0.002 (statistical) in MC. There is an additional 1.9% relative systematic uncertainty between the data and MC, bringing the total uncertainty to 2.6%. Further adding uncertainties from the material assay brings this to 3.2% absolute uncertainty. The data response is 3% higher than the MC predicts, a little more than the total relative uncertainty. The MC predicts a resolution of 11.5%, which is an adequate description of the low statistics data.

The MC predicts the response in the TE configuration is 3% higher, because most electrons ionize their way through the tracker before electromagnetic showers develop in the ECAL. This sample provides another 62 events in this sample, more positrons than electrons because of the running conditions. A similar 3% discrepancy response is seen in these TE results. The systematics are somewhat larger but the statistical uncertainty is smaller, because of slightly better statistics and resolution. The data/MC relative total uncertainty is 2.4%, and the absolute uncertainty on the response is 3.1%. The MC prediction of a 9.1% resolution describes the data well.

This sample is subject to the same systematics as the proton and pion results plus additional 1.1% uncertainty due to the electron-pure selection. Comparing the default MC to a variation with ±1.2% Pb density in the ECAL reveals only a ±0.15% change in response for the TE configuration and ±0.3% change for the EH sample. Variations of the event selection contributes 1% uncertainty to the response. The absolute energy scale uncertainty is the same 2% and the data vs. MC relative uncertainty is 0.7% from the material model effects and calibrations described previously. Another 0.5% comes from the cross talk model. Finally, the beam momentum uncertainty is 1% at these energies.

For some neutrino analyses, especially with electron or neutral pion final states, the accuracy of the electromagnetic response is of interest. For high energy hadronic showers, it is traditional to form the e/pi response ratio. A shower initiated by a charged hadron will typically have both hadronic and electromagnetic components, the relative fractions of each evolve with energy

and vary stochastically event by event. An e/pi response calibration illustrates the accuracy of these components separately. In this detector, the comparison of the higher 0.8 for electrons to the lower 0.6 response for pions provides a helpful rule of thumb.

10. Calorimetry discussion

In addition to the classic and exhaustive monograph on calorimetry by Richard Wigmans [17], there are several recent test beam measurements using segmented scintillator and absorber calorimeters and hadron simulations similar to ours.

The MINOS neutrino experiment uses a detector made of scintillator and inch-thick iron, very similar to the MINERvA HCAL. Their test beam exposure in the CERN T7 and T11 beamlines was analyzed to produce electron [18] and hadron [19] calorimetry results. They compared their data to a GEANT3 simulation and found several kinds of discrepancies at the 3% to 6% level. However, our data are compared to a Geant4 simulation, so comparison to the present analysis is difficult.

The CALICE experiment has operated several kinds of sampling calorimeters in beams at Fermilab and CERN. They use similar, Geant4 based hadron and electromagnetic models, but their data is mostly at higher energy. The analysis of their data is ongoing. As of this writing, two publications [20, 21] make for good comparison to the MINERvA test beam data.

Hadronic calorimetry is considered in [21] with data taken with an iron-scintillator calorimeter. They find Geant4 models underestimate the measured response by 2% at 8 GeV/c momentum, their lowest pion data available. This is also the only data point in their analysis where Geant4 is using the same Bertini Cascade model used in our simulation. Their data show the trend switches and the MC overestimates the data above 20 GeV or so, but remain consistent within their uncertainty estimates.

In the later paper [20], data from a Tungsten segmented ECAL calorimeter is

compared to Geant4 models for electrons, pions, and protons. The π^+ response for the same Bertini cascade model (but from Geant4 9.6.p2) describes their mean response very well from 3 to 8 GeV. The discrepancy is less than 2% while their uncertainty is around 3%. Agreement also follows for proton data in the same range. A similar result is obtained for positrons, agreement above 2 GeV. However, the simulation underestimates the data by 2.5% at 1 GeV, just within 1σ agreement for the lowest positron energy for which they have data.

Taken together, the MINERvA and CALICE data suggest that the Bertini cascade model from recent (9.4p2 and later) Geant4 does a good job of describing hadronic data at the 3% level in an iron-scintillator calorimeter through the combined range of energy. CALICE indicates that the electromagnetic cascade model applied to an ECAL style calorimeter also does very good. But the low energy data point that is similar to MINERvA's suggests the MC underestimates the response in both cases.

In general, hadron calorimetry at energies below 2 GeV follows a process where one hadron undergoes zero or one inelastic interaction, with a small number of outgoing particles. It is not easily characterized by the statistical $\sqrt{E}$ and $\sqrt{N}$ effects typical of hadron calorimetry at higher energies.

11. Tracking validation

The proton sample in the TE detector configuration allows us to validate proton tracking efficiency: the probability that a proton will be reconstructed as a three-dimensional track object. The proton tracking efficiency, and that for pions, is important for measurements of neutrino differential cross sections with specific proton and pion final states.

The sample is similar to the one used for the Birks' parameter measurement where protons stop no later than plane 19, but without the requirement that its depth be consistent with a proton at the end of its range. Another difference, the sample is extended downstream to protons whose last activity is plane 6. This tests a combination of the standard MINERvA "long tracker" which requires a

minimum of 11 planes and two variations of the short tracker which can form tracks with as few as five planes. For this analysis, the MC sample is four times the size of the data sample.

For this sample, protons with kinetic energy less than 0.4 GeV whose last energy deposit is between planes 9 and 19 (inclusive) are tracked with efficiency of $99.2^{+0.2}_{-0.3}$% in data and 99.8±0.1% in MC. For the data, this corresponds to tracking 1520 out of 1533 protons in the sample. Around 60% of protons stopping a distance consistent with the end of their range, and failing the tracking is highly correlated with a proton experiencing an interaction.

Differences begin to appear for even shorter proton samples. For the 185 protons that appear to stop in plane 8, 178 of them were tracked, which gives $96.2^{+1.2}_{-1.6}$% compared to the MC $97.7^{+0.5}_{-0.6}$%. For 338 protons that appear to stop in planes 6 and 7 only 308 are tracked, $91.1^{+1.5}_{-1.6}$% compared to the MC 96.5±0.5%. These subsamples have a 70% fraction with their stopping location at the end of their expected range. It is more likely in the data that a short event at the end of its expected range will not pass the tracking requirements.

The above results for protons were obtained with a short tracker configured for a neutrino pion production analysis [5]. A somewhat different configuration optimized for a quasi-elastic proton analysis [6] gives 1 to 2% higher efficiency, successfully tracking an additional 6, 1, and 8 events in the data subsamples for the shortest, 8-plane, and longest samples respectively, with a similar trend of better tracking in the MC. The results from the test beam data are then applied to both proton and pion tracking in the neutrino analyses.

The main reason for the difference between the two tracking techniques involves the choice of candidate clusters of activity to include in the tracking. The QE proton style analysis is more permissive, especially allowing clusters with more hits and more energy that would be expected from a simply ionizing particle. The pion analysis excludes these when deciding whether to form a track. In the case of very short, six-plane tracks, excluding one plane has a large effect.

These results suggest that tracking efficiency is adequately modeled (within

1%) for tracks greater than 9 planes, which makes it a negligible uncertainty for neutrino analyses. In contrast, we can use a data-based correction of as much as 5% to the efficiency for shorter track lengths, relative to the MC predicted efficiency. The tracking algorithm also has to deal with activity near the neutrino interaction point and wider range of angles relative to the detector axis, but not in the test beam sample. Therefore, this efficiency correction should be on top of the MC prediction for efficiency considering all effects seen in real neutrino interactions.

12. Conclusion

We have measured the performance of the tracking and calorimetry of the MINERvA detector design by exposing a miniature version of the detector to a test beam of low momentum protons, pions, and electrons from the Fermilab Test Beam Facility. These data provide a constraint on the Birks' law saturation effect for our formulation of polystyrene based plastic scintillator. The calorimetric response to protons and pions within the range of energies tested yields uncertainties of 3% when the single particle calorimetric response is used in neutrino analyses. The electron sample yields a similar uncertainty. Tracking performance is well modeled, and we have measured a small discrepancy between the performance of tracking in the data and simulation. In summary, there are several effects that could be interpreted as 2σ fluctuations relative to the systematic uncertainties, but overall the MC describes the data and its resolutions well.

References

[1] L. Aliaga, et al., Design, Calibration, and Performance of the MINERvA Detector, Nucl.Instrum.Meth. A743 (2014) 130–159. arXiv:1305.5199, doi:10.1016/j.nima.2013.12.053.

[2] L. Fields, J. Chvojka, et al., Measurement of Muon Antineutrino Quasi-Elastic Scattering on a Hydrocarbon Target at $E_\nu \sim 3.5$ GeV, Phys. Rev.

Lett. 111, 022501. arXiv:1305.2234, doi:10.1103/PhysRevLett.111.022501.

[3] G. Fiorentini, D. Schmitz, P. Rodriguez, et al., Measurement of Muon Neutrino Quasi-Elastic Scattering on a Hydrocarbon Target at $E_\nu \sim 3.5$ GeV, Phys. Rev. Lett. 111, 022502. arXiv:1305.2243, doi:10.1103/PhysRevLett.111.022502.

[4] B. Tice, M. Datta, J. Mousseau, et al., Measurement Ratios of ν_μ Charged-Current Cross Sections on C, Fe, and Pb to CH at Neutrino Energies 2-20 GeV, Phys. Rev. Lett. 112, 231801. arXiv:1403.2103, doi:10.1103/PhysRevLett.112.231801.

[5] B. Eberly, et al., Charged Pion Production in ν_μ Interactions on Hydrocarbon at $\langle E_\nu \rangle = 4.0$ GeV arXiv:1406.6415.

[6] T. Walton, M. Betancourt, et al., Measurement of muon plus proton final states in ν_μ Interactions on Hydrocarbon at $< E_\nu > = 4.2$ GeV arXiv:1409.4497.

[7] N. Tagg, et al., Arachne - A web-based event viewer for MINERvA, Nucl.Instrum.Meth. 676 (2012) 44–49. arXiv:1111.5315, doi:10.1016/j.nima.2012.01.059.

[8] S. Agostinelli, et al., GEANT4: A Simulation toolkit, Nucl.Instrum.Meth. A506 (2003) 250–303. doi:10.1016/S0168-9002(03)01368-8.

[9] J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Dubois, et al., Geant4 developments and applications, IEEE Trans.Nucl.Sci. 53 (2006) 270. doi:10.1109/TNS.2006.869826.

[10] J. B. Birks, , Proc.Phys.Soc. A64 (1951) 874.

[11] J. B. Birks, The Theory and Practice of Scintillation Counting.

[12] V. Tretyak, Semi-empirical calculation of quenching factors for ions in scintillators, Astropart.Phys. 33 (2010) 40–53. arXiv:0911.3041, doi:10.1016/j.astropartphys.2009.11.002.

[13] L. Reichhart, D. Y. Akimov, H. Araujo, E. Barnes, V. Belov, et al., Quenching Factor for Low Energy Nuclear Recoils in a Plastic Scintillator, Phys.Rev. C85 (2012) 065801. arXiv:1111.2248, doi:10.1103/PhysRevC.85.065801.

[14] A. Heikkinen, N. Stepanov, J. P. Wellisch, Bertini intranuclear cascade implementation in GEANT4, eConf C0303241 (2003) MOMT008. arXiv:nucl-th/0306008.

[15] D. Ashery, I. Navon, G. Azuelos, H. Walter, H. Pfeiffer, et al., True Absorption and Scattering of Pions on Nuclei, Phys.Rev. C23 (1981) 2173–2185. doi:10.1103/PhysRevC.23.2173.

[16] B. Allardyce, C. Batty, D. Baugh, E. Friedman, G. Heymann, et al., Pion reaction cross-sections and nuclear sizes, Nucl.Phys. A209 (1973) 1–51. doi:10.1016/0375-9474(73)90049-3.

[17] R. Wigmans, Calorimetry: Energy Measurement in Particle Physics, Int.Ser.Monogr.Phys. 107 (2000) 1–726.

[18] P. L. Vahle, Electromagnetic interactions in the MINOS detectors (Ph.D thesis 2004).

[19] M. A. Kordosky, Hadronic interactions in the MINOS detectors (Ph.D. thesis 2004).

[20] C. Adloff, et al., Validation of GEANT4 Monte Carlo Models with a Highly Granular Scintillator-Steel Hadron Calorimeter, JINST 8 (2013) 07005. arXiv:1306.3037, doi:10.1088/1748-0221/8/07/P07005.

[21] C. Adloff, J. J. Blaising, M. Chefdeville, C. Drancourt, R. Gaglione, et al., Shower development of particles with momenta from 1 to 10 GeV in the

CALICE Scintillator-Tungsten HCAL, JINST 9 (2014) P01004. arXiv:1311.3505, doi:10.1088/1748-0221/9/01/P01004.